arXiv:1401.1269v1 [stat.CO] 7 Jan 2014

Bayesian Robust Inference of Sample Selection Using Selection-t Models

Peng Ding
Department of Statistics, Harvard University
Address: One Oxford Street, Cambridge, MA 02138, USA
E-mail: pengding@fas.harvard.edu

Abstract

Heckman selection model is the most popular econometric model in analysis of data with sample selection. However, selection models with Normal errors cannot accommodate heavy tails in the error distribution. Recently, Marchenko and Genton proposed a selection-t model to perform frequentist' robust analysis of sample selection. Instead of using their maximum likelihood estimates, our paper develops new Bayesian procedures for the selection-t models with either continuous or binary outcomes. By exploiting the Normal mixture representation of the t distribution, we can use data augmentation to impute the missing data, and use parameter expansion to sample the restricted covariance matrices. The Bayesian procedures only involve simple steps, without calculating analytical or numerical derivatives of the complicated log likelihood functions. Simulation studies show the vulnerability of the selection models with Normal errors, as well as the robustness of the selection models with t errors. Interestingly, we find evidence of heavy-tailedness in three real examples analyzed by previous studies, and the conclusions about the existence of selection effect are very sensitive to the distributional assumptions of the error terms.

Key Words: Data augmentation; Heavy-tailedness; Parameter expansion; Restricted covariance matrix; Sample selection.

1 Introduction

Sample selection often occurs in social sciences and biomedical studies, when the outcomes of interest are partially observed or the samples are not representative of the population. In the analysis of labor market, Heckman (1979) proposed a selection model comprising of a Probit sample selection equation and a Normal linear outcome equation. The sample selection problem arises, when the error terms of the sample selection equation and the outcome equation are correlated. Heckman (1979) treated the sample selection as a model misspecification problem due to an omitted variable, and proposed a two-step procedure to adjust the linear regression model with an extra nonlinear term. The Heckman selection model, or Type II Tobit model, is now widely used in many fields. Heckman's two-step procedure and the maximum likelihood estimation (MLE) using Newton-Raphson iteration are incorporated into many standard econometrical and statistical routines, such as the `heckman` procedure in `Stata` (StataCorp 2013) and the `sampleSelection` (Toomet and Henningsen 2008) package of `R` (R Development Core Team 2010).

Despite its popularity and wide applications, researchers tried to generalize the Heckman selection model in various ways. Examples include transformation-based model (Lee 1983), semiparametric model (Ahn and Powell 1993), and nonparametric model (Das, Newey and Vella 2003). Among the Bayesian community, Li (1998) proposed a Bayesian inference procedure for the Heckman selection model using data augmentation (Tanner and Wong 1987), Chib, Greenberg and Jeliazkov (2009) generalized it to Bayesian nonparametric model, and Van Hasselt (2011) generalized it to Bayesian semiparametric model using Dirichlet process prior.

Models based on t distributions are widely applied for robust analysis (Albert and Chib 1993; Geweke 1992; Liu 1999; Liu 2004), and they are attractive alternatives for the models based on Normal distribution such as linear and Probit models. Recently, Marchenko and Genton (2012) extended the selection model to deal with heavy-tailedness by modeling the error terms as a bivariate t distribution, which was called a selection-t model. Marchenko and Genton (2012) proposed a Newton-Raphson iteration

procedure to find the MLE of the selection-t model. However, the frequentists' solution has several limitations: first, it involves complicated derivatives of the log likelihood function; second, it is not very direct to be generalized to binary outcomes; third, the inference based on the asymptotic Normality of MLE may not be accurate in problems with small sample sizes. In order to overcome these limitations, we propose Bayesian procedures for the selection-t model using Markov Chain Monte Carlo (MCMC). The Bayesian procedures exploit the Normal mixture representation of the t distribution (Albert and Chib 1993), and use data augmentation (Tanner and Wong 1987) to impute the latent variables. However, difficulty arises when sampling the covariance matrix of the error terms. For the purpose of full identification, the variance of the error term in the selection equation is restricted to be one, which makes the posterior distribution of the covariance matrix non-standard and difficult to sample directly. Previous studies (Koop and Poirier 1997; Li, 1998; McCulloch, Polson and Rossi 2000; Van Hasselt 2011) reparametrized the restricted covariance matrix. However, we use parameter expansion (Liu and Rubin 1998; Meng and Van Dyk 1999; Van Dyk and Meng 2001; Imai and Van Dyk 2005) to overcome the sampling difficulty of the restricted covariance matrix, which allows us to have conjugate distributions. We provide a more detailed discussion of the two classes of approaches in Section 8.

The remainder of the article proceeds as follows. In Section 2, we briefly review the Heckman selection model. Section 3 introduces the selection-t model (Marchenko and Genton 2012) and its Normal mixture representation, which is fundamental to our Bayesian procedures. Section 4 presents a Bayesian inference procedure for the selection-t model, and Section 5 generalizes it to deal with binary outcome. We show some simulation studies to evaluate the finite sample properties of our procedures in Section 6. In Section 7, we apply our new Bayesian procedures to empirical applications, where we find evidence of heavy-tailedness. Section 8 contains some discussions and possible extensions. All the technical details are shown in Appendix. Throughout this article, all vectors are column vectors, and we use boldface letters to represent vectors and matrices. Functions written in `R` for the methods proposed in this paper are available upon request from the author.

2 Review of the Heckman Selection Model

A selection model has two parts: a regression equation for the outcome, and a regression equation for the sample selection mechanism. Suppose the regression equation for the outcome of primary interest is

$$y_i^* = \boldsymbol{x}_i^\top \boldsymbol{\beta} + \varepsilon_i, \tag{1}$$

and the sample selection mechanism is driven by the following latent linear regression equation

$$u_i^* = \boldsymbol{w}_i^\top \boldsymbol{\gamma} + \eta_i, \tag{2}$$

for $i = 1, \cdots, N$. The covariates in $\boldsymbol{x}_i$ and $\boldsymbol{w}_i$ may overlap with each other, and the exclusion restriction holds when at least one of the elements of $\boldsymbol{w}_i$ are not in $\boldsymbol{x}_i$. Let K and L denote the dimensions of $\boldsymbol{x}_i$ and $\boldsymbol{w}_i$, respectively. We observe the outcome y_i^*, if and only if $u_i^* > 0$. Therefore, the indicator for sample selection is

$$u_i = I(u_i^* > 0). \tag{3}$$

Let y_i be the observed outcome, with $y_i = y_i^*$ if $u_i = 1$, and $y_i = \text{NA}$ is $u_i = 0$, where "NA" indicates missing data.

Heckman (1979) assumed a bivariate Normal distribution for ε_i and η_i:

$$\begin{pmatrix} \varepsilon_i \\ \eta_i \end{pmatrix} \sim \boldsymbol{N}_2(\boldsymbol{0}_2, \boldsymbol{\Omega}), \tag{4}$$

where $\boldsymbol{0}_2 = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$ and $\boldsymbol{\Omega} = \begin{pmatrix} \sigma_1^2 & \rho\sigma_1 \\ \rho\sigma_1 & 1 \end{pmatrix}$. We fix the second diagonal element of $\boldsymbol{\Omega}$ at 1 for full identification. Under the bivariate Normal assumption, the mean equation for the outcomes of the selected samples is

$$E\{y \mid u = 1, \boldsymbol{x}, \boldsymbol{w}\} = \boldsymbol{x}^\top \boldsymbol{\beta} + \rho\sigma_1 \lambda(\boldsymbol{w}^\top \boldsymbol{\gamma}), \tag{5}$$

where $\lambda(\cdot) = \phi(\cdot)/\Phi(\cdot)$ is the inverse Mills ratio. Therefore, the sample selection problem can be treated as a model misspecification problem, because the mean equation for the outcomes of the selected samples is a linear function $\boldsymbol{x}_i^\top \boldsymbol{\beta}$ with a nonlinear

correction term $\rho\sigma_1\lambda(\boldsymbol{w}^\top\boldsymbol{\gamma})$. Based on (5), Heckman (1979) proposed a two-step procedure by first fitting a Probit model of u on $\boldsymbol{w}$ to obtain $\widehat{\boldsymbol{\gamma}}$, and then fitting a linear model of y on $\left\{\boldsymbol{x}, \lambda(\boldsymbol{w}^\top\widehat{\boldsymbol{\gamma}})\right\}$ to obtain $(\widehat{\boldsymbol{\beta}}, \widehat{\rho}, \widehat{\sigma}_1)$. The two-step procedure is less efficient than the full information MLE, but it is robust to the deviation of the joint Normality of the error terms. The MLE of the Heckman selection model can be calculated by Newton-Raphson iteration or EM algorithm (Little and Rubin 2002). Alternatively, Bayesian posterior inference of the Heckman selection model can be obtained by data augmentation (Li 1998).

3 Normal Mixture Representation of the Selection-t Model

In order to model heavy-tailedness, Marchenko and Genton (2012) proposed a selection-t model, and assumed that ε_i and η_i follow a bivariate t distribution with unknown degrees of freedom ν, namely,

$$\begin{pmatrix}\varepsilon_i \\ \eta_i\end{pmatrix} \sim \boldsymbol{t}_2(\boldsymbol{0}_2, \boldsymbol{\Omega}, \nu). \tag{6}$$

The density function of the bivariate t distribution $\boldsymbol{t}_2(\boldsymbol{\mu}, \boldsymbol{\Omega}, \nu)$ is

$$f(\boldsymbol{t}; \boldsymbol{\mu}, \boldsymbol{\Omega}, \nu) = (2\pi)^{-1}|\boldsymbol{\Omega}|^{-1/2}\left\{1 + \nu^{-1}(\boldsymbol{t}-\boldsymbol{\mu})^\top\boldsymbol{\Omega}^{-1}(\boldsymbol{t}-\boldsymbol{\mu})\right\}^{-(\nu+2)/2}. \tag{7}$$

As $\nu \to +\infty$, the bivariate t distribution in (6) converges to the bivariate Normal distribution in (4). Thus, the Heckman selection model is a limiting case of the selection-t model. In this article, we use the name "selection model" for (4) and the name "selection-t model" for (6).

However, the density of the t distribution in (7) results in cumbersome posterior distributions, which can be solved by using data augmentation. By introducing latent variables $\{q_i : i = 1, \cdots, N\}$, the bivariate t distribution of ε_i and η_i has the following Normal mixture representation:

$$\begin{pmatrix}\varepsilon_i \\ \eta_i\end{pmatrix} \sim \boldsymbol{N}_2(\boldsymbol{0}_2, \alpha\boldsymbol{\Omega}/q_i), \text{ where } q_i \sim \alpha\chi^2_\nu/\nu, i = 1, \cdots, N. \tag{8}$$

The parameter α is not identifiable from the observed data $\boldsymbol{D}_{obs} = \{(y_i, u_i, \boldsymbol{x}_i, \boldsymbol{w}_i) : i = 1, \cdots, N\}$, but it is identifiable from the complete data $\boldsymbol{D}_{com} = \{(y_i^*, u_i^*, y_i, u_i, q_i, \boldsymbol{x}_i, \boldsymbol{w}_i) :$

$i = 1, \cdots, N\}$. When α is fixed at one, the model is fully identifiable. The over-parametrization for this model is a way of parameter expansion to accelerate the convergence rates of the MCMC samplers (Liu and Wu 1999; Meng and Van Dyk 1999; Van Dyk and Meng 2001).

For Bayesian inference, we need to specify prior distributions for all the parameters $(\boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\Omega}, \nu, \alpha)$. We choose a multivariate Normal prior for the regression coefficients $(\boldsymbol{\beta}^\top, \boldsymbol{\gamma}^\top)^\top \sim \boldsymbol{N}_{K+L}(\boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0)$, a Gamma prior for the degrees of freedom, $\nu \sim \text{Gamma}(\alpha_0, \beta_0)$ with a shape parameter α_0 and a rate parameter β_0, and a scaled-inverse-χ^2 prior for α, $\alpha \sim b/\chi^2_c$.

We restrict the second diagonal element of the covariance matrix $\boldsymbol{\Omega}$ to be one, which makes it difficult to sample $\boldsymbol{\Omega}$ from its posterior distribution directly. In order to use parameter expansion, we consider the unrestricted covariance matrix $\boldsymbol{\Sigma} = \text{diag}\{1, \sigma_2\}\ \boldsymbol{\Omega}\ \text{diag}\{1, \sigma_2\}$. The Inverse-Wishart prior $\boldsymbol{W}_2^{-1}(\nu_0, \boldsymbol{I}_2)$ for the covariance matrix $\boldsymbol{\Sigma}$ is equivalent to the priors for $(\boldsymbol{\Omega}, \sigma_2^2)$:

$$f(\boldsymbol{\Omega}) \propto (1-\rho^2)^{-3/2}\sigma_1^{-(\nu_0+3)} \exp\left\{-\frac{1}{2\sigma_1^2(1-\rho^2)}\right\}, \tag{9}$$

$$\text{and } \sigma_2^2|\boldsymbol{\Omega} \sim \{(1-\rho^2)\chi^2_{\nu_0}\}^{-1}. \tag{10}$$

Details of the derivations for (9) and (10) are in Appendix A.

4 Bayesian Computation for the Selection-t Model

Bayesian computation using data augmentation includes two main steps: the imputation step (I-step) by imputing the missing data $\boldsymbol{D}_{mis} = \{(y_i^*, u_i^*, q_i) : i = 1, \cdots, N\}$, and the posterior step (P-step) by updating the posterior distributions of the parameters.

For the ease of derivation, we introduce the following matrix notation. The joint model of the latent outcome and the selection mechanism is

$$\begin{pmatrix} y_i^* \\ u_i^* \end{pmatrix} = \begin{pmatrix} \boldsymbol{x}_i^\top & \boldsymbol{0}_L \\ \boldsymbol{0}_K & \boldsymbol{w}_i^\top \end{pmatrix} \begin{pmatrix} \boldsymbol{\beta} \\ \boldsymbol{\gamma} \end{pmatrix} + \begin{pmatrix} \varepsilon_i \\ \eta_i \end{pmatrix}.$$

Define $\boldsymbol{Z}_i = \begin{pmatrix} y_i^* \\ u_i^* \end{pmatrix}$ as the latent outcome and selection mechanism, $\boldsymbol{V}_i = \begin{pmatrix} \boldsymbol{x}_i^\top & \boldsymbol{0}_L \\ \boldsymbol{0}_K & \boldsymbol{w}_i^\top \end{pmatrix}$ as the design matrix of the covariates, and $\boldsymbol{\delta} = \begin{pmatrix} \boldsymbol{\beta} \\ \boldsymbol{\gamma} \end{pmatrix}$ as the regression coefficients. The

complete data likelihood is

$$
\begin{aligned}
& \alpha^{-n}|\boldsymbol{\Omega}|^{-n/2}\exp\left\{-\frac{1}{2\alpha}\sum_{i=1}^{N} q_i(\boldsymbol{Z}_i-\boldsymbol{V}_i\boldsymbol{\delta})^\top\boldsymbol{\Omega}^{-1}(\boldsymbol{Z}_i-\boldsymbol{V}_i\boldsymbol{\delta})\right\} \\
& \cdot\prod_{i=1}^{n} q_i\frac{1}{(2\alpha/\nu)^{\nu/2}\Gamma(\nu/2)}q_i^{\nu/2-1}e^{-\nu q_i/(2\alpha)} \\
\propto\ & \alpha^{-n-n\nu/2}(\nu/2)^{n\nu/2}\{\Gamma(\nu/2)\}^{-n}|\boldsymbol{\Omega}|^{-n/2}\exp\left\{-\frac{1}{2\alpha}\sum_{i=1}^{N} q_i(\boldsymbol{Z}_i-\boldsymbol{V}_i\boldsymbol{\delta})^\top\boldsymbol{\Omega}^{-1}(\boldsymbol{Z}_i-\boldsymbol{V}_i\boldsymbol{\delta})\right\} \\
& \cdot\prod_{i=1}^{n} q_i^{\nu/2}e^{-\nu q_i/(2\alpha)}. \qquad (11)
\end{aligned}
$$

4.1 The Imputation Step

First, we impute the missing data given the observed data and the parameters. Let $TN(\mu,\sigma^2;L,U)$ be a Normal distribution $N(\mu,\sigma^2)$ truncated within the interval $[L,U]$. In the imputation step, we sample $(\alpha,\{\boldsymbol{Z}_i,q_i\}|\{y_i,u_i\},\boldsymbol{\delta},\boldsymbol{\Omega},\nu)$ jointly, and then marginalize over α by discarding its sample. Since $\alpha|(\{y_i,u_i\},\boldsymbol{\delta},\boldsymbol{\Omega},\nu)$ is the same as its prior distribution, we draw $\alpha\sim b/\chi^2_c$. Given $(u_i^*,y_i,u_i,\boldsymbol{\delta},\boldsymbol{\Omega},\nu,\alpha)$, we use Gibbs sampler to draw $\{\boldsymbol{Z}_i\}$ and $\{q_i\}$ iteratively. The conditional means and variances of the bivariate Normal variables (y_i^*,u_i^*) have the following forms:

$$\mu_{u|y}=\boldsymbol{w}_i^\top\boldsymbol{\gamma}+\rho(y_i^*-\boldsymbol{x}_i^\top\boldsymbol{\beta})/\sigma_1,\quad \sigma^2_{u|y}=\alpha(1-\rho^2)/q_i, \qquad (12)$$

$$\mu_{y|u}=\boldsymbol{x}_i^\top\boldsymbol{\beta}+\rho\sigma_1(u_i^*-\boldsymbol{w}_i^\top\boldsymbol{\gamma}),\quad \sigma^2_{y|u}=\alpha\sigma_1^2(1-\rho^2)/q_i. \qquad (13)$$

Given $(q_i,y_i,u_i,\boldsymbol{\delta},\boldsymbol{\Omega},\nu,\alpha)$, we impute the latent variables (y_i^*,u_i^*) as follows: if $u_i=1$, we draw

$$y_i^*=y_i, \qquad (14)$$

$$\text{and}\qquad u_i^*|\,(y_i^*,q_i,y_i,u_i,\boldsymbol{\delta},\boldsymbol{\Omega},\nu,\alpha)\sim TN(\mu_{u|y},\sigma^2_{u|y};0,\infty); \qquad (15)$$

if $u_i=0$, we draw

$$u_i^*|\,(q_i,y_i,u_i,\boldsymbol{\delta},\boldsymbol{\Omega},\nu,\alpha)\sim TN(\boldsymbol{w}_i^\top\boldsymbol{\gamma},\alpha/q_i;-\infty,0), \qquad (16)$$

$$\text{and}\qquad y_i^*|\,(u_i^*,q_i,y_i,u_i,\boldsymbol{\delta},\boldsymbol{\Omega},\nu,\alpha)\sim N(\mu_{y|u},\sigma^2_{y|u}). \qquad (17)$$

Given $(\boldsymbol{Z}_i, y_i, u_i, \boldsymbol{\delta}, \boldsymbol{\Omega}, \nu, \alpha)$, we impute q_i by the scaled-inverse-χ^2 distribution:

$$q_i | \left(\boldsymbol{Z}_i, y_i, u_i, \boldsymbol{\delta}, \boldsymbol{\Omega}, \nu, \alpha\right) \sim \alpha \chi^2_{\nu+2} / \left\{(\boldsymbol{Z}_i - \boldsymbol{V}_i\boldsymbol{\delta})^\top \boldsymbol{\Omega}^{-1} (\boldsymbol{Z}_i - \boldsymbol{V}_i\boldsymbol{\delta}) + \nu\right\}. \tag{18}$$

Algorithm 1 summarizes the imputation step for the selection-t model.

Algorithm 1 Imputation step for the selection-t model

I-1 Draw $\alpha \sim b/\chi^2_c$ from its prior distribution;
I-2 Draw (y_i^*, u_i^*) according to (14) to (17) for $i = 1, \cdots, N$;
I-3 Draw q_i according to (18) for $i = 1, \cdots, N$.

4.2 Posterior Step

Second, we draw the parameters from their posterior distributions, conditioning on the complete data $\boldsymbol{D}_{com}$. After imputing the missing data, the parameter α is identifiable from the complete data, and it follows a scaled-inverse-χ^2 distribution:

$$\alpha | \left(\{\boldsymbol{Z}_i, q_i, \boldsymbol{V}_i\}, \boldsymbol{\delta}, \boldsymbol{\Omega}, \nu\right) \sim \left[b + \sum_{i=1}^{N} q_i \{(\boldsymbol{Z}_i - \boldsymbol{V}_i\boldsymbol{\delta})^\top \boldsymbol{\Omega}^{-1} (\boldsymbol{Z}_i - \boldsymbol{V}_i\boldsymbol{\delta}) + \nu\}\right] \Big/ \chi^2_{c+2N+N\nu}. \tag{19}$$

The complete data likelihood in (11) demonstrates the Normality of $\boldsymbol{\delta} | \left(\{\boldsymbol{Z}_i, \boldsymbol{V}_i\}, \boldsymbol{\Omega}, \nu, \alpha\right)$, because of the quadratic log posterior density. The posterior mean and precision matrix of $\boldsymbol{\delta}$ are determined by the mode and the negative Hessian matrix of the log posterior density. Since the log of the conditional posterior density of $\boldsymbol{\delta}$ is

$$-\frac{1}{2\alpha} \sum_{i=1}^{N} q_i (\boldsymbol{Z}_i - \boldsymbol{V}_i\boldsymbol{\delta})^\top \boldsymbol{\Omega}^{-1} (\boldsymbol{Z}_i - \boldsymbol{V}_i\boldsymbol{\delta}) - \frac{1}{2} (\boldsymbol{\delta} - \boldsymbol{\mu}_0)^\top \boldsymbol{\Sigma}_0^{-1} (\boldsymbol{\delta} - \boldsymbol{\mu}_0),$$

we draw

$$\boldsymbol{\delta} | \left(\{\boldsymbol{Z}_i, \boldsymbol{V}_i\}, \boldsymbol{\Omega}, \nu, \alpha\right) \sim \boldsymbol{N}_{K+L}(\widehat{\boldsymbol{\mu}}_{\boldsymbol{\delta}}, \widehat{\boldsymbol{\Sigma}}_{\boldsymbol{\delta}}), \tag{20}$$

where

$$\widehat{\boldsymbol{\mu}}_{\boldsymbol{\delta}} = \widehat{\boldsymbol{\Sigma}}_{\boldsymbol{\delta}} \left(\sum_{i=1}^{N} q_i \boldsymbol{V}_i^\top \boldsymbol{\Omega}^{-1} \boldsymbol{Z}_i / \alpha + \boldsymbol{\Sigma}_0^{-1} \boldsymbol{\mu}_0\right) \text{ and } \widehat{\boldsymbol{\Sigma}}_{\boldsymbol{\delta}} = \left(\sum_{i=1}^{N} q_i \boldsymbol{V}_i^\top \boldsymbol{\Omega}^{-1} \boldsymbol{V}_i / \alpha + \boldsymbol{\Sigma}_0^{-1}\right)^{-1}.$$

Given $(\{\boldsymbol{Z}_i, \boldsymbol{V}_i, q_i\}, \boldsymbol{\delta}, \nu, \alpha)$, it is difficult to sample the restricted covariance matrix $\boldsymbol{\Omega}$ directly. However, parameter expansion allows us to reparametrize the model and get conjugate posterior distributions. Define

$$\boldsymbol{E}_i = \begin{pmatrix} 1 & 0 \\ 0 & \sigma_2 \end{pmatrix} (\boldsymbol{Z}_i - \boldsymbol{V}_i\boldsymbol{\delta}), \tag{21}$$

and we have $\boldsymbol{E}_i|\,(q_i,\boldsymbol{\delta},\nu,\alpha) \sim \boldsymbol{N}_2(\boldsymbol{0}_2,\alpha\boldsymbol{\Sigma}/q_i)$. Since the prior $\boldsymbol{\Sigma} \sim \boldsymbol{W}_2^{-1}(\nu_0,\boldsymbol{I}_2)$ implies the priors in (9) and (10), we first draw $\sigma_2^2|\boldsymbol{\Omega} \sim \{(1-\rho^2)\chi^2_{\nu_0}\}^{-1}$, and then transform the data to get $\boldsymbol{E}_i$ using (21). The conditional posterior of $\boldsymbol{\Sigma}$ is

$$\begin{aligned}
& |\boldsymbol{\Sigma}|^{-(\nu_0+3)/2}\exp\left\{-\frac{1}{2}\mathrm{tr}(\boldsymbol{\Sigma}^{-1})\right\}\prod_{i=1}^{N}|\alpha\boldsymbol{\Sigma}/q_i|^{-1/2}\exp\left\{-\frac{1}{2\alpha}\sum_{i=1}^{N}q_i\boldsymbol{E}_i^\top\boldsymbol{\Sigma}^{-1}\boldsymbol{E}_i\right\} \\
\propto\ & |\boldsymbol{\Sigma}|^{-(N+\nu_0+3)/2}\exp\left\{-\frac{1}{2}\mathrm{tr}\left[\boldsymbol{\Sigma}^{-1}(\boldsymbol{S}+\boldsymbol{I}_2)\right]\right\} \\
\sim\ & \boldsymbol{W}_2^{-1}(N+\nu_0,\boldsymbol{S}+\boldsymbol{I}_2),
\end{aligned}$$

where $\boldsymbol{S}=\sum_{i=1}^{N}q_i\boldsymbol{E}_i\boldsymbol{E}_i^\top/\alpha$. Therefore, we draw $\boldsymbol{\Sigma}|\,(\{\boldsymbol{E}_i,q_i\},\boldsymbol{\delta},\nu,\alpha) \sim \boldsymbol{W}_2^{-1}(N+\nu_0,\boldsymbol{S}+\boldsymbol{I}_2)$, and transform $\boldsymbol{\Sigma}$ to

$$\sigma_2^2=\sigma_{22} \text{ and } \boldsymbol{\Omega}=\begin{pmatrix}1 & 0\\ 0 & 1/\sigma_2\end{pmatrix}\boldsymbol{\Sigma}\begin{pmatrix}1 & 0\\ 0 & 1/\sigma_2\end{pmatrix}. \tag{22}$$

Given $(\{\boldsymbol{Z}_i,\boldsymbol{V}_i,q_i\},\boldsymbol{\delta},\boldsymbol{\Omega},\alpha)$, the conditional posterior density of ν is:

$$f(\nu|\{\boldsymbol{Z}_i,\boldsymbol{V}_i,q_i\},\boldsymbol{\delta},\boldsymbol{\Omega},\alpha) \propto \exp\left\{N\nu\log(\nu/2)/2-N\log\Gamma(\nu/2)+(\alpha_0-1)\log\nu-\xi\nu\right\}, \tag{23}$$

where

$$\xi=\beta_0+N\log\alpha/2+\sum_{i=1}^{N}q_i/(2\alpha)-\sum_{i=1}^{N}\log q_i/2.$$

Unfortunately, the conditional distribution of ν is not standard. Geweke (1992) proposed a rejection sampling method using an exponential distribution as a proposal density. Albert and Chib (1993) were interested in the posterior probabilities for ν in a finite set, and they suggested sampling ν from a discrete distribution. In our studies, we treat ν as a continuous parameter as Geweke (1992). We advocate a more accurate Gamma approximation, and its shape parameter α^* and rate parameter β^* are discussed in Appendix C. The approximate Gamma distribution is a proposal density for the Metropolized Independence Sampler (Liu 2001) for ν, which is a special case of the Metropolis-Hastings Algorithm. In our simulation studies and real examples, the Gamma approximation works fairly well with acceptance rates higher than 0.95. According to Liu (2001), the efficiency of the Metropolized Independence Sampler depends on how close the proposal density is to the target density. Therefore, the optimal acceptance rates of the Metropolized Independence Sampler are usually higher than the random walk Metropolis algorithm.

Algorithm 2 summarizes the posterior step for the selection-t model.

Algorithm 2 Posterior step for the selection-t model

P-1 Draw α according to (19);

P-2 Draw $\boldsymbol{\delta} \sim \boldsymbol{N}_{K+L}(\widehat{\boldsymbol{\mu}}_{\boldsymbol{\delta}}, \widehat{\boldsymbol{\Sigma}}_{\boldsymbol{\delta}})$ according to (20);

P-3 Draw σ_2^2 from its prior $\{(1-\rho^2)\chi^2_{\nu_0}\}^{-1}$, make the transformation (21), draw $\boldsymbol{\Sigma} \sim \boldsymbol{W}_2^{-1}(N+\nu_0, \boldsymbol{S}+\boldsymbol{I}_2)$, and transform to $\boldsymbol{\Omega}$ according to (22);

P-4 Given the old value ν, draw a proposal $\nu' \sim \text{Gamma}(\alpha^*, \beta^*)$, with acceptance probability

$$\min\left\{1, \frac{f(\nu'|\cdot)/\text{dgamma}(\nu', \alpha^*, \beta^*)}{f(\nu|\cdot)/\text{dgamma}(\nu, \alpha^*, \beta^*)}\right\},$$

where $f(\nu|\cdot)$ is the density defined in (23), and $\text{dgamma}(\nu, \alpha^*, \beta^*)$ is the Gamma density evaluated at ν.

5 The Selection-Robit Model

In this section, we will discuss the selection-t model with binary outcomes, which can be easily generalized to other types of limited dependent outcomes. We assume that the regression models for the outcome and the selection mechanism are the same as (1) and (2), but the observed outcome is $y_i = I(y_i^* > 0)$ if $u_i = 1$, and $y_i = NA$ if $u_i = 0$. In order to get full identification, we assume that ε_i and η_i follow

$$\begin{pmatrix}\varepsilon_i \\ \eta_i\end{pmatrix} \sim \boldsymbol{t}_2(\boldsymbol{0}_2, \boldsymbol{R}, \nu),$$

with a Normal mixture representation

$$\begin{pmatrix}\varepsilon_i \\ \eta_i\end{pmatrix} \sim \boldsymbol{N}_2(\boldsymbol{0}_2, \alpha\boldsymbol{R}/q_i), \text{ where } q_i \sim \alpha\chi^2_\nu/\nu, i = 1, \cdots, N; \text{ and } \boldsymbol{R} = \begin{pmatrix}1 & \rho \\ \rho & 1\end{pmatrix}.$$

Since we can only observe the signs of the latent outcome and selection mechanism, the variances of the error terms are restricted to be one for full identification.

A regression model with binary outcomes can be represented as a latent linear model, with only the signs of the latent outcomes observed (Albert and Chib 1993). Different distributions of the error terms correspond to different generalized linear models, including Normal (Probit model), logistic (Logit model) and t (Robit model; Liu 2004). Therefore, we call the selection model with binary outcomes "selection-Probit" model and selection-t model with binary outcomes "selection-Robit" model.

For the selection-Robit model, direct sampling the correlation matrix $\boldsymbol{R}$ involves non-standard distributions. Again, we solve this problem by parameter expansion. The Inverse-Wishart $\boldsymbol{W}_2^{-1}(\nu_0, \boldsymbol{I}_2)$ prior for $\boldsymbol{\Sigma} = \text{diag}\{\sigma_1, \sigma_2\}\ \boldsymbol{R}\ \text{diag}\{\sigma_1, \sigma_2\}$ is equivalent to the priors for $(\boldsymbol{R}, \sigma_1^2, \sigma_2^2)$:

$$f(\rho) \propto (1-\rho^2)^{-(\nu_0-3)/2}, \tag{24}$$

$$\text{and } \sigma_j^2|\rho \sim \{(1-\rho^2)\chi^2_{\nu_0}\}^{-1} (j = 1, 2). \tag{25}$$

Details of the derivations of (24) of (25) are in Appendix B. If we choose $\nu_0 = 3$, the prior distribution for ρ is Uniform$(-1, 1)$. Therefore, the prior $\boldsymbol{W}_2^{-1}(3, \boldsymbol{I}_2)$ for $\boldsymbol{\Sigma}$ is a marginally uniform prior (Barnard, McCulloch and Meng 2000).

For selection-Robit model, the imputation of $\boldsymbol{Z}_i$ changes slightly. The conditional means and variances of u_i^* and y_i^* have the following forms:

$$\tilde{\mu}_{u|y} = \boldsymbol{w}_i^\top \boldsymbol{\gamma} + \rho(y_i^* - \boldsymbol{x}_i^\top \boldsymbol{\beta}), \quad \tilde{\sigma}^2_{u|y} = \alpha(1-\rho^2)/q_i,$$
$$\tilde{\mu}_{y|u} = \boldsymbol{x}_i^\top \boldsymbol{\beta} + \rho(u_i^* - \boldsymbol{w}_i^\top \boldsymbol{\gamma}), \quad \tilde{\sigma}^2_{y|u} = \alpha(1-\rho^2)/q_i.$$

If $u_i = 1$, we draw the truncated bivariate Normal distribution using the Gibbs sampler:

$$u_i^* | \left(y_i^*, q_i, y_i, u_i, \boldsymbol{\delta}, \rho, \nu, \alpha\right) \sim TN(\tilde{\mu}_{u|y}, \tilde{\sigma}^2_{u|y}; 0, \infty), \tag{26}$$

$$y_i^* | \left(u_i^*, q_i, y_i, u_i, \boldsymbol{\delta}, \rho, \nu, \alpha\right) \sim y_i TN(\tilde{\mu}_{y|u}, \tilde{\sigma}^2_{y|u}; 0, \infty) + (1-y_i) TN(\tilde{\mu}_{y|u}, \tilde{\sigma}^2_{y|u}; -\infty, 0). \tag{27}$$

If $u_i = 0$ and y_i is missing, we impute the missing data by two steps:

$$u_i^* | \left(q_i, y_i, u_i, \boldsymbol{\delta}, \rho, \nu, \alpha\right) \sim TN(\boldsymbol{w}_i^\top \boldsymbol{\gamma}, \alpha/q_i; -\infty, 0), \tag{28}$$

$$y_i^* | \left(u_i^*, q_i, y_i, u_i, \boldsymbol{\delta}, \rho, \nu, \alpha\right) \sim N(\tilde{\mu}_{y|u}, \tilde{\sigma}^2_{y|u}). \tag{29}$$

The full conditional distributions of $\boldsymbol{\delta}$ and ν have the same forms as in the previous section, except that $\boldsymbol{\Omega}$ is replaced by $\boldsymbol{R}$. The full conditional distribution of $\boldsymbol{R}$ does not have conjugate form, which can be circumvented by parameter expansion. Define

$$\tilde{\boldsymbol{E}}_i = \begin{pmatrix} \sigma_1 & 0 \\ 0 & \sigma_2 \end{pmatrix} (\boldsymbol{Z}_i - \boldsymbol{V}_i \boldsymbol{\delta}), \tag{30}$$

and we have $\tilde{\boldsymbol{E}}_i|\,(q_i,\boldsymbol{\delta},\nu,\alpha) \sim \boldsymbol{N}_2(\boldsymbol{0}_2,\alpha\boldsymbol{\Sigma}/q_i)$. The prior $\boldsymbol{\Sigma} \sim \boldsymbol{W}_2^{-1}(\nu_0,\boldsymbol{I}_2)$ implies prior for $(\boldsymbol{R},\sigma_1^2,\sigma_2^2)$ as shown in (24) and (25). We first independently draw $\sigma_j^2|\boldsymbol{\Omega} \sim \{(1-\rho^2)\chi^2_{\nu_0}\}^{-1}(j=1,2)$, and then transform the data to obtain $\tilde{\boldsymbol{E}}_i$ using (30). The conditional posterior of $\boldsymbol{\Sigma}$ is $\boldsymbol{W}_2^{-1}(N+\nu_0,\tilde{\boldsymbol{S}}+\boldsymbol{I}_2)$, where $\tilde{\boldsymbol{S}}=\sum_{i=1}^N q_i\tilde{\boldsymbol{E}}_i\tilde{\boldsymbol{E}}_i^\top/\alpha$. Therefore, we draw $\boldsymbol{\Sigma}|(\{\tilde{\boldsymbol{E}}_i,q_i\},\boldsymbol{\delta},\nu,\alpha) \sim \boldsymbol{W}_2^{-1}(N+\nu_0,\tilde{\boldsymbol{S}}+\boldsymbol{I}_2)$, and transform $\boldsymbol{\Sigma}$ to

$$\sigma_1^2=\sigma_{11},\sigma_2^2=\sigma_{22} \text{ and } \boldsymbol{R}=\begin{pmatrix}1/\sigma_1 & 0\\ 0 & 1/\sigma_2\end{pmatrix}\Sigma\begin{pmatrix}1/\sigma_1 & 0\\ 0 & 1/\sigma_2\end{pmatrix}. \tag{31}$$

Algorithm 3 and 4 summarize the imputation and posterior steps for the selection-Robit model, respectively.

Algorithm 3 Imputation step for the selection-Robit model

I′-1 Draw $\alpha \sim b/\chi_c^2$;
I′-2 Impute (y_i^*,u_i^*) according to (26) to (27);
I′-3 Draw q_i the same as I-3 except that $\boldsymbol{\Omega}$ is replaced by $\boldsymbol{R}$.

Algorithm 4 Imputation step for the selection-Robit model

P′-1 Draw α the same as P-1 except that $\boldsymbol{\Omega}$ is replaced by $\boldsymbol{R}$;
P′-2 Draw $\boldsymbol{\delta}$ the same as P-2 except that $\boldsymbol{\Omega}$ is replaced by $\boldsymbol{R}$;
P′-3 Draw σ_1^2 and σ_2^2 independently from their priors $\{(1-\rho^2)\chi^2_{\nu_0}\}^{-1}$, make the transformation (30), draw $\boldsymbol{\Sigma} \sim \boldsymbol{W}_2^{-1}(N+\nu_0,\tilde{\boldsymbol{S}}+\boldsymbol{I}_2)$, and transform $\boldsymbol{\Sigma}$ back to $\boldsymbol{R}$ according to (31);
P′-4 Draw ν the same as P-4.

6 Simulation Studies

In order to evaluate the finite sample properties of the new Bayesian procedures, several simulation studies are presented in this section. The MCMC algorithms 1 to 4 for the selection-t model and selection-Robit model can be easily refined to be the algorithms for the selection model and selection-Probit model, by restricting $\alpha=1$, $\nu=+\infty$ and $q_i=1(i=1,\cdots,N)$. Throughout our simulation studies and our empirical studies, the parameters for the prior distributions are chosen as follows: $\boldsymbol{\mu}_0=\boldsymbol{0}_{K+L},\boldsymbol{\Sigma}_0=\text{diag}\{1,\cdots,1\}/100,\nu_0=3,\alpha_0=1,\beta_0=0.1,b=0.1$, and $c=0.1$. The results are not sensitive to other choices of hyperparameters. We run the MCMC algorithms for 5×10^4 iterations, with the first 10^4 draws discarded as a burn in period. The results

from multiple chains differ very slightly and all of them converge with Gelman-Rubin diagnostic statistics close to 1, and we only present the result from a single chain in each of our simulation study and real application.

6.1 Data Generated from Models with Normal Errors

Selection-t models are more general than the selection models, and they converge to the selection models when $\nu \to +\infty$. We first generate the observed data from a selection model with bivariate Normal errors. We generate the covariates from $x_i \sim N(0, 2^2), w_i \sim N(0, 2^2)$, and x_i is independent of w_i; and generate the latent outcome and selection mechanism from $y_i^* = 0.5 + x_i + \varepsilon_i, u_i^* = 2 + x_i + 1.5w_i + \eta_i$, and

$$\begin{pmatrix} \varepsilon_i \\ \eta_i \end{pmatrix} \sim \boldsymbol{N}_2 \left[\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \boldsymbol{\Omega}_0 = \begin{pmatrix} 1 & 0.3 \\ 0.3 & 1 \end{pmatrix} \right].$$

The selection indicator is $u_i = I(u_i^* > 0)$, and the outcome is $y_i = y_i^*$ if $u_i = 1$ and $y_i = \text{NA}$ if $u_i = 0$. In our generated data set, about 30% outcomes are missing. We apply the Bayesian procedures for both the selection model and selection-t model, and the posterior distributions of the parameters are summarized in the boxplots in Figure 1(a). The boxplots in white and grey are obtained under the selection model and selection-t model, respectively. The posterior distributions of the parameters $(\beta_1, \gamma_1, \gamma_2, \rho)$ concentrate near their true values under both models. And the posterior draws of ν under the selection-t model take very large values. It is known that t distributions with large degrees of freedom approximate the Normal distribution. Therefore, large values of the posterior draws of ν are evidence of Normality.

Replacing y_i with $I(y_i > 0)$ in the same data set, we implement the Bayesian procedures for both the selection-Probit and the selection-Robit model. The posterior distributions of the parameters are summarized in the boxplots in Figure 1(b), with the boxplots in white for the selection-Probit model and the boxplots in grey for the selection-Robit model. Although more diffused than the distributions in Figure 1(a) due to loss of information after dichotomizing the continuous outcomes, the posterior distributions of the parameters in Figure 1(b) concentrate near the true parameters. And the large values of posterior draws of ν under the selection-Robit model show evidence of Normality.

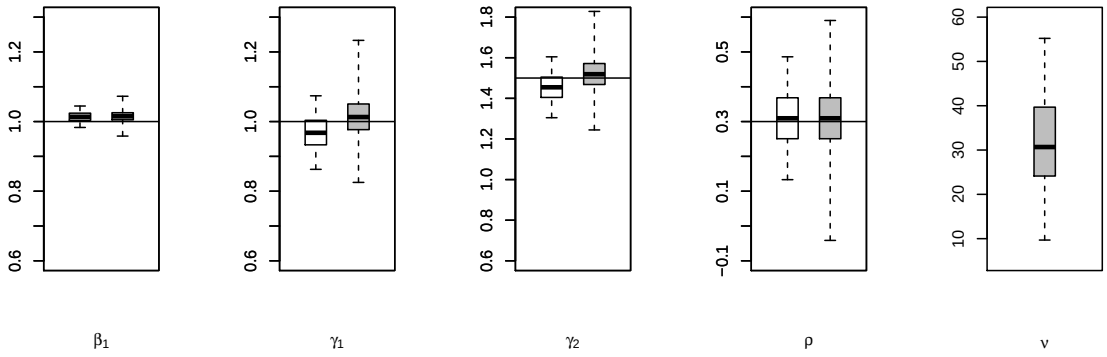

(a) Data generated from the selection model with Normal errors and analyzed by the selection model (white) and the Selection-t model (grey).

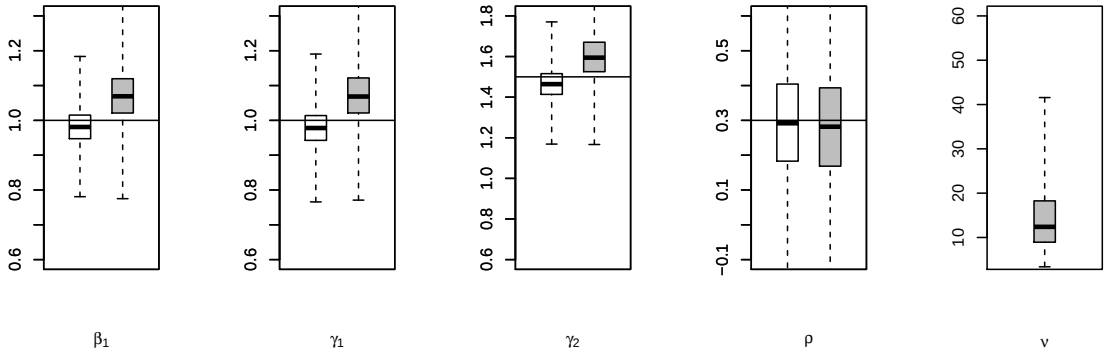

(b) Data generated from the Selection-Probit model and analyzed by Selection-Probit model (white) and Selection-Robit model (grey).

Figure 1: Selection Models with Normal Errors

6.2 Data Generated from Models with t Errors

We generate data from a selection-t model. The data generating process is the same as Section 6.1 except that the bivariate Normal distribution is replaced by a bivariate t distribution with degrees of freedom $\nu = 3$. In our generated data set, about 30% outcomes are missing. The posterior distributions of the parameters are summarized in the boxplots in Figure 2(a) under both the selection model (in white) and the selection-t model (in grey). The posterior distributions under the selection-t model are close to the true parameters, but those under the selection model are far from the true parameters (e.g., γ_1 and γ_2 in Figure 2(a)). Again, replacing y_i with $I(y_i > 0)$, we apply both the

Bayesian procedures for the selection-Probit model and the selection-Robit model, and the boxplots of the posterior distributions are shown in Figure 2(b). We can see that the posterior distributions of $(\beta_1, \gamma_1, \gamma_2)$ under the selection-Probit model are farther from the true parameters than those under the selection-Robit model.

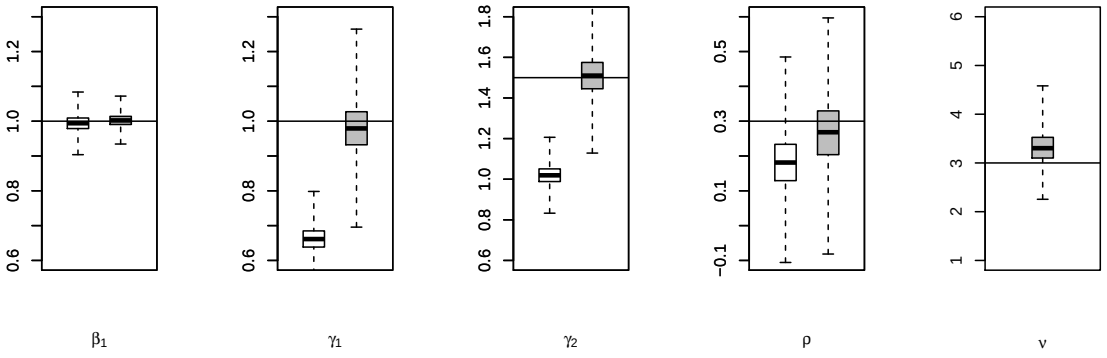

(a) Data generated from the Selection-t model with Normal errors and analyzed by the selection model (white) and the Selection-t model (grey).

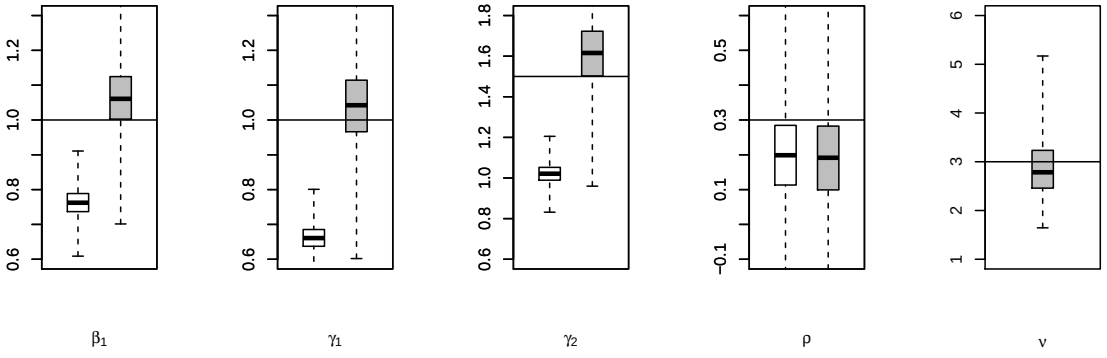

(b) Data generated from the Selection-Robit model and analyzed by Selection-Probit model (white) and Selection-Robit model (grey).

Figure 2: Selection Models with t Errors

6.3 Data Generated from Models with Gaussian Mixture Errors

We generate data from a selection model with Gaussian Mixture errors:

$$0.4\boldsymbol{N}_2(\mathbf{0}_2, \boldsymbol{\Omega}_0) + 0.3\boldsymbol{N}_2(\mathbf{0}_2, 2\boldsymbol{\Omega}_0) + 0.2\boldsymbol{N}_2(\mathbf{0}_2, 4\boldsymbol{\Omega}_0) + 0.1\boldsymbol{N}_2(\mathbf{0}_2, 8\boldsymbol{\Omega}_0) + 0.1\boldsymbol{N}_2(\mathbf{0}_2, 16\boldsymbol{\Omega}_0),$$

where $\mathbf{\Omega}_0$ is defined in Section 6.1. In our generated data set, about 33% of the outcomes are missing. Figure 3(a) are the boxplots of the posterior distributions of the parameters under the selection model (in white) and selection-t model (in grey), and Figure 3(b) are the boxplots of the posterior distributions of the parameters under the selection-Probit (in white) and selection-Robit model (in grey). Since both the selection models based on Normal and t distribution are misspecified, most posterior distributions do not concentrate very near the true parameters. The behaviors of the selection and selection-Probit models are very wild (e.g. γ_1 and γ_2) due to the heavy-tails generated by the Gaussian Mixture errors, while the behaviors of the selection-t and selection-Robit model are much more robust to the model misspecifications.

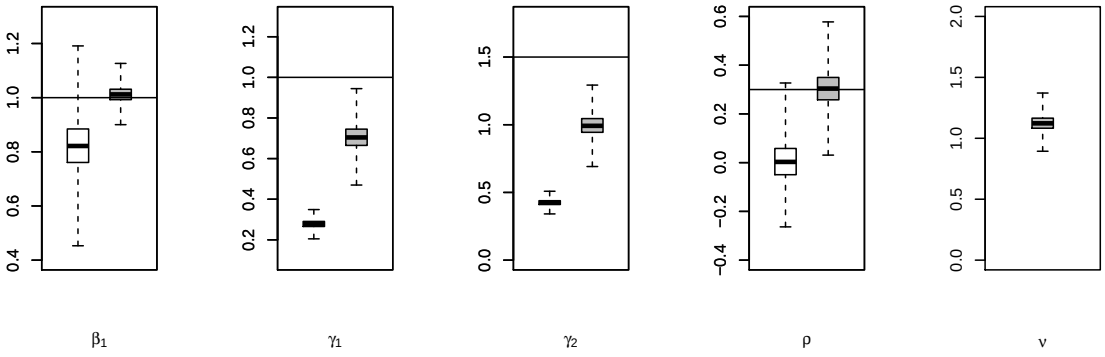

(a) Data generated from the Selection model with Gaussian Mixture errors and analyzed by the selection model (white) and the Selection-t model (grey).

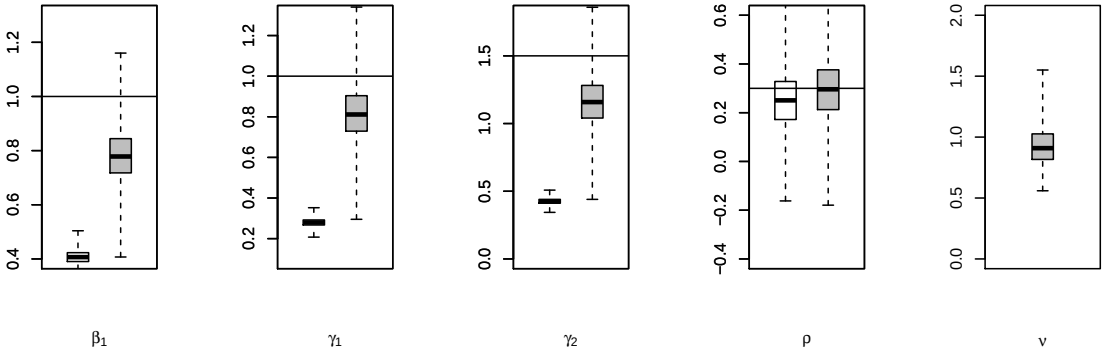

(b) Data generated from the Selection model with Gaussian Mixture errors and analyzed by Selection-Probit model (white) and Selection-Robit model (grey).

Figure 3: Selection Models with Gaussian Mixture Errors

7 Empirical Studies

7.1 Ambulatory Expenditures

We apply our Bayesian procedures to the data about ambulatory expenditures. The data are taken from Cameron and Trivedi (2010), which were re-analyzed by Marchenko and Genton (2012) using frequentists' procedure for the selection-t model. The data can be downloaded from `http://cameron.econ.ucdavis.edu/musbook/mus.html`. In our analysis, we choose log expenditures (`lambexp`) as the outcome variable y. The covariates in the outcome equation are $\boldsymbol{x} = (1, \texttt{age}, \texttt{female}, \texttt{educ}, \texttt{blhisp}, \texttt{totchr}, \texttt{ins})$, including age, gender, education status, ethnicity, number of chronic diseases and insurance status. The exclusion restriction assumption holds by including the income variable into the selection equation, i.e., $\boldsymbol{w} = (\boldsymbol{x}, \texttt{income})$. In order to compare with the frequentists' approach, we choose the same set of covariates as Marchenko and Genton (2012). We would expect similar results from both frequentists' and Bayesian procedures, because the prior information will be overwhelmed by the data with relatively large sample size ($n = 3328$ with 526 missing outcomes). Table 1 shows that the results of frequentists' and Bayesian methods under the same model are very close to each other, while the selection model and the selection-t model give us different results. The consistent results from both the frequentists' procedure and the Bayesian procedure verify our MCMC code and convergence of the Markov Chains indirectly.

The value of ρ measures the sample selection effect, with $\rho = 0$ indicating the absence of the sample selection bias. Under the selection model, both the frequentists' 95% confidence interval and the Bayesian 95% posterior credible interval of ρ contain zero, which indicates weak evidence of the sample selection bias. However, under the selection-t model, neither the frequentists' 95% confidence interval nor the Bayesian 95% posterior credible interval of ρ contains zero, which suggests the existence of sample selection effect. Different inferences on the existence of selection effect give us different statistical and economic interpretations. Absence of selection effect ($\rho = 0$) implies that the outcomes are missing at random (Little and Rubin, 2002), and the observed outcomes are representative for inference of the ambulatory expenditures given the

Table 1: The Ambulatory Expenditures Example

	Selection		Selection-t		Selection (Bayes)		Selection-t (Bayes)	
Outcome Model								
age	0.212	$(0.167, 0.257)$	0.207	$(0.163, 0.251)$	0.211	$(0.165, 0.256)$	0.207	$(0.162, 0.251)$
female	0.348	$(0.230, 0.466)$	0.307	$(0.196, 0.417)$	0.339	$(0.217, 0.456)$	0.306	$(0.194, 0.417)$
educ	0.019	$(-0.002, 0.039)$	0.017	$(-0.003, 0.037)$	0.018	$(-0.003, 0.037)$	0.017	$(-0.003, 0.038)$
blhisp	-0.219	$(-0.336, -0.102)$	-0.193	$(-0.306, -0.080)$	-0.213	$(-0.329, -0.094)$	-0.193	$(-0.306, -0.078)$
totchr	0.540	$(0.463, 0.617)$	0.513	$(0.443, 0.583)$	0.534	$(0.453, 0.611)$	0.512	$(0.441, 0.583)$
ins	-0.030	$(-0.130, 0.070)$	-0.053	$(-0.151, 0.046)$	-0.033	$(-0.133, 0.068)$	-0.054	$(-0.153, 0.046)$
Selection Model								
age	0.088	$(0.034, 0.142)$	0.099	$(0.040, 0.157)$	0.088	$(0.033, 0.142)$	0.099	$(0.040, 0.157)$
female	0.663	$(0.543, 0.782)$	0.725	$(0.591, 0.859)$	0.664	$(0.544, 0.784)$	0.729	$(0.597, 0.867)$
educ	0.062	$(0.038, 0.086)$	0.065	$(0.040, 0.090)$	0.062	$(0.038, 0.085)$	0.065	$(0.040, 0.090)$
blhisp	-0.364	$(-0.485, -0.243)$	-0.394	$(-0.524, -0.263)$	-0.364	$(-0.485, -0.244)$	-0.394	$(-0.525, -0.265)$
totchr	0.797	$(0.658, 0.936)$	0.890	$(0.719, 1.061)$	0.795	$(0.660, 0.936)$	0.893	$(0.733, 1.075)$
ins	0.170	$(0.047, 0.293)$	0.180	$(0.048, 0.313)$	0.169	$(0.045, 0.291)$	0.180	$(0.047, 0.314)$
income	0.003	$(0.000, 0.005)$	0.003	$(0.000, 0.006)$	0.003	$(0.000, 0.005)$	0.003	$(0.000, 0.006)$
σ	1.271	$(1.236, 1.308)$	1.195	$(1.146, 1.246)$	1.277	$(1.241, 1.324)$	1.195	$(1.148, 1.249)$
ρ	-0.131	$(-0.401, 0.161)$	-0.322	$(-0.526, -0.083)$	-0.159	$(-0.462, 0.108)$	-0.327	$(-0.536, -0.085)$
ν	$+\infty$		12.938	$(8.391, 19.917)$	$+\infty$		12.913	$(8.841, 22.447)$

Note: MLEs and 95% confidence intervals of the selection model (column 2 and 3), the selection-t model (column 4 and 5); Bayesian posterior means and 95% credible intervals of the selection model (column 6 and 7), the selection-t model (column 8 and 9).

observed covariates. The results of the selection model and the selection-t model are different, because of the heavy-tailedness in the data. Figure 4(a) depicts the posterior distribution of ν with 95% posterior credible interval $(8.841, 22.447)$, which is evidence of heavy-tailedness.

7.2 Wage Offer Function for Married Women

We re-analyze the data from Mroz (1987) and Wooldridge (2002) to estimate the wage offer function for married women. The data set can be found in the `R` package `sampleSelection`. The outcome of interest is the log of `wage`, which are missing for 325 individuals and observed for 428 individuals. The covariates in the outcome equation are $\boldsymbol{x} = (1, \texttt{educ}, \texttt{exper}, \texttt{exper}^2)$, including education status, experience and its squared term. The covariates in the selection equation includes other income, age, number of young children and number of older children as additional variables, i.e., $\boldsymbol{w} = (\boldsymbol{x}, \texttt{nwifeinc}, \texttt{age}, \texttt{kids5}, \texttt{kids618})$. Wooldridge (2002) used Heckman's two-step procedure, because the MLE in this particular data is numerically unstable. The `heckman` procedure in old versions of `Stata` (e.g. `Stata` 10) converges to a local maxi-

mum. The results in Table 2 are from the latest version of `Stata` (`Stata` 13), which are the same as the results from `R` function `selection()` in the package `sampleSelection` (`R` version 2.10.0 with `sampleSelection` package version 0.7.2). Bayesian procedures are more preferable in this case, because they are less sensitive to the initial values. Table 2 shows that the Bayesian and frequentists' methods for the selection model give us similar results. Most regression coefficients under the selection model and the selection-t model do not differ dramatically. However, there does exist some different economic interpretations under different models. For instance, according to the results from Bayesian inference, the `exper` variable enter the outcome equation as $-0.001\texttt{exper}^2 + 0.043\texttt{exper}$ and $0.026\texttt{exper}$ under the selection model and selection-t model, respectively. The former implies that log(`wage`) is a quadratic function of `exper` with a maximum point at 21.5, while the latter implies that it is a linear thus increasing function of `exper`.

Table 2: The Wage Offer Example

	Selection		Selection-t		Selection (Bayes)		Selection-t (Bayes)	
Outcome Model								
educ	0.108	$(0.079, 0.137)$	0.108	$(0.085, 0.132)$	0.108	$(0.078, 0.137)$	0.109	$(0.085, 0.132)$
exper	0.043	$(0.014, 0.072)$	0.025	$(0.001, 0.050)$	0.043	$(0.012, 0.072)$	0.026	$(0.001, 0.051)$
exper^2	-0.001	$(-0.002, 0.000)$	-0.000	$(-0.001, 0.002)$	-0.001	$(-0.002, 0.000)$	-0.000	$(-0.001, 0.002)$
Selection Model								
educ	0.131	$(0.082, 0.181)$	0.140	$(0.079, 0.201)$	0.132	$(0.083, 0.182)$	0.143	$(0.083, 0.208)$
exper	0.123	$(0.087, 0.160)$	0.150	$(0.104, 0.195)$	0.124	$(0.087, 0.161)$	0.151	$(0.105, 0.199)$
exper^2	-0.002	$(-0.003, -0.001)$	-0.002	$(-0.004, -0.001)$	-0.002	$(-0.003, -0.001)$	-0.002	$(-0.004, -0.001)$
nwifeinc	-0.012	$(-0.022, -0.003)$	-0.012	$(-0.024, 0.000)$	-0.012	$(-0.022, -0.003)$	-0.012	$(-0.024, -0.001)$
age	-0.053	$(-0.069, -0.036)$	-0.064	$(-0.085, -0.043)$	-0.053	$(-0.069, -0.036)$	-0.065	$(-0.087, -0.045)$
kids5	-0.867	$(-1.100, -0.635)$	-1.043	$(-1.339, -0.749)$	-0.870	$(-1.101, -0.640)$	-1.060	$(-1.374, -0.771)$
kids618	0.036	$(-0.049, 0.121)$	0.034	$(-0.070, 0.139)$	0.037	$(-0.048, 0.123)$	0.034	$(-0.071, 0.142)$
σ	0.663	$(0.619, 0.708)$	0.444	$(0.394, 0.500)$	0.670	$(0.627, 0.720)$	0.451	$(0.399, 0.508)$
ρ	0.027	$(-0.262, 0.315)$	-0.383	$(-0.632, -0.061)$	0.019	$(-0.308, 0.281)$	-0.362	$(-0.605, -0.053)$
ν	$+\infty$		3.061	$(2.340, 4.183)$	$+\infty$		3.094	$(2.296, 4.318)$

Note: MLEs and 95% confidence intervals of the selection model (column 2 and 3), the selection-t model (column 4 and 5); Bayesian posterior means and 95% credible intervals of the selection model (column 6 and 7), the selection-t model (column 8 and 9).

Under the selection model, the 95% confidence interval and posterior credible interval of ρ cover zero, and the evidence for the sample selection effect is weak. However, the result using the selection-t model differs a lot, because the posterior credible interval of ρ does not cover zero, and it indicates the existence of the sample selection

effect. The conclusion changes, because ν is very small with posterior credible interval $(2.296, 4.318)$ and histogram in Figure 4(b). The tail behavior of t distribution with small degrees of freedom is very different from the tail behavior of the normal distribution. From the summary statistics of the observed values of the outcome, we find that the kurtosis of the observed outcome is greater than five, which is much greater than three, the kurtosis of normal distributions. Heavy-tailedness problem in this data is very severe, and the selection-t model is more preferable than the selection model.

7.3 HIV Survey with Nonparticipation

We apply our Bayesian procedure for selection-Robit model to the data from the 2007 Zambia Demographic and Health Survey (Bärnighausen et al., 2011). We are interested in estimating the HIV prevalence rate among men. However, in the survey, 6416 male participants were tested for HIV status, but 1318 males rejected to take the HIV test. For those nonparticipants of the HIV test, their binary outcomes y_i (HIV status) are missing with $u_i = 0$. The covariates in the outcome equation are $\boldsymbol{x} = (1, \texttt{age}, \texttt{income}, \texttt{marital}, \texttt{condom}, \texttt{evertestHIV}, \texttt{smoking}, \texttt{location})$, including age, income, marital status, condom use, ever tested HIV, smoking status, and the location of the individuals. And the covariates in the selection equation are $\boldsymbol{w}$ including both $\boldsymbol{x}$ and the identity of the interviewer, with the latter included for the exclusion restriction assumption. In Table 3, we show the MLEs and 95% confidence intervals from the `heckprob` procedure in `Stata`, and the Bayesian posterior medians and 95% credible intervals of the selection-Robit model. The Bayesian posterior distributions of the selection-Probit is omitted here, since the extreme draws of ρ from the Gibbs sampler make the algorithm numerically unstable. Therefore, we find the numerical stability in presence of heavy-tailedness in the outcome as another advantage of the selection-Robit model. The frequentists' procedure for the selection-Robit is also omitted, since it is not very straightforward to modify Marchenko and Genton (2012)'s procedure to deal with binary outcomes.

Due to the extremely heavy tails in the data ($\nu < 1$ for most posterior draws in Figure 4(c)), the regression coeffficients under selection-Probit and selection-Robit

Table 3: The HIV Survey Example

	Selection-Probit (`heckprob`)		Selection-Robit (Bayes)	
Outcome Model				
age	0.035	$(0.007, 0.065)$	0.037	$(-0.016, 0.095)$
income	-0.157	$(-0.308, -0.007)$	-0.355	$(-0.772, 0.070)$
marital	0.708	$(0.517, 0.898)$	2.179	$(1.407, 4.595)$
condom	0.266	$(0.130, 0.402)$	0.715	$(0.448, 1.011)$
evertestHIV	0.167	$(0.058, 0.277)$	0.297	$(0.063, 0.534)$
smoking	0.002	$(-0.113, 0.117)$	-0.005	$(-0.271, 0.248)$
Selection Model				
age	0.024	$(0.001, 0.046)$	0.303	$(0.045, 0.484)$
income	-0.159	$(-0.270, -0.047)$	-0.762	$(-1.489, 0.512)$
marital	-0.074	$(-0.176, 0.029)$	-1.001	$(-1.941, 0.255)$
condom	0.055	$(-0.041, 0.151)$	-0.141	$(-1.402, 0.672)$
evertestHIV	0.045	$(-0.040, 0.130)$	0.540	$(-0.135, 1.096)$
smoking	0.160	$(0.072, 0.248)$	1.043	$(-0.268, 1.884)$
ρ	-0.590	$(-0.863, 0.125)$	-0.968	$(-0.999, -0.737)$
ν	∞		0.378	$(0.330, 0.487)$

Note: MLEs and 95% confidence intervals of the selection-Probit model (column 2 and 3); Bayesian posterior medians and 95% credible intervals of the selection-Robit model (column 4 and 5).

models differ a lot. And the conclusion about selection effect is much more significant under the selection-Robit model than under the selection-Probit model. The presence of selection effect makes more practical sense in this particular example. Since the HIV positive individuals tend to hide their status and reject to participate in the HIV test, the error terms ε_i and η_i should be negatively correlated.

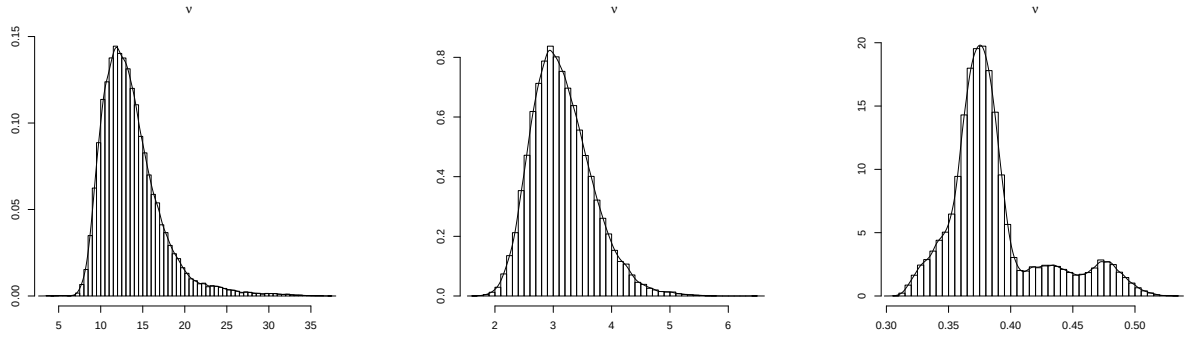

(a) Ambulatory expenditures (b) Wage offer (c) HIV survey

Figure 4: Posterior Distribution of ν in Three Empirical Studies

8 Discussion

This paper develops new Bayesian procedures for a class of selection-t models, which was recently proposed by Marchenko and Genton (2012) to deal with heavy-tailedness of the data. Selection models with t errors are robust parametric alternatives for the selection models with Normal errors (Heckman 1979). Although selection-t models are not as flexible as semiparametric (Chib et al. 2009) and nonparametric (Van Hasselt 2011) models, they can model data with heavy tails by introducing only one extra parameter ν for controlling the heavy-tailedness (Marchenko and Genton 2012). Efficient implementations of the selection models are realized by Bayesian procedures using data augmentation and parameter expansion. We illustrate the potential applications of our Bayesian procedures for selection-t models with three real problems. The heavy-tailedness seems very common in practice, since we find strong evidence of heavy-tailedness in all of our empirical studies. In our examples, the conclusions about the existence of the selection effect differ dramatically under different distributional assumptions of the error terms, and results from different models may provide different practical interpretations. It is our future research direction to study the generalized selection-t models with different degrees of freedom for the outcome equation and selection equation.

Due to the identification issue, the variances of the error terms in the selection equation in the selection and selection-t models are restricted to be one, and the variances of the error terms in both the selection and outcome equations in the selection-Probit and selection-Robit models are restricted to be one. These restrictions make the posterior distributions of the covariance matrix nonstandard and nonconjugate, which complicate the MCMC procedures. One possible solution of this problem is to reparametrize the covariance matrix (Koop and Poirier 1997; Li 1998; McCulloch, Polson and Rossi 2000; Van Hasselt 2011) in terms of the covariance $\text{cov}(\varepsilon_i, \eta_i) = \sigma_{12}$ and the conditional variance $\text{var}(\varepsilon_i \mid \eta_i) = \sigma^2_{1|2}$:

$$\mathbf{\Omega} = \begin{pmatrix} \sigma_1^2 & \rho\sigma_1 \\ \rho\sigma_1 & 1 \end{pmatrix} = \begin{pmatrix} \sigma^2_{1|2} + \sigma^2_{12} & \sigma_{12} \\ \sigma_{12} & 1 \end{pmatrix}.$$

Imposing inverse-Gamma prior on $\sigma^2_{1|2}$ and Normal prior on σ_{12} will result in conjugate

conditional posterior distributions of the parameters. Our paper takes an alternative perspective, and the MCMC algorithms rely on the parameter expansion technique by enlarging the parameter space to maintain conjugacy, which "is closely related to reparametrization techniques" (Liu and Wu 1999). Thus, the previous reparametrization method and our parameter expansion method share some common essence. In analysis of the multinomial Probit model, Imai and Van Dyk (2005) used the parameter expansion method, and showed some advantages of the parameter expansion approach in terms of convergence rate of the MCMC algorithms. Although our algorithms require imputing all the missing y_i^* which may adversely affect the mixing properties, the parameter expansion seems to be a compensation of this drawback in terms of convergence rate. Our future research will focus on combination the strength of both approaches. Although reparametrization in Koop and Poirier (1997), Li (1998), McCulloch, Polson and Rossi (2000) and Van Hasselt (2011) is very convenient to deal with selection models, the implied prior for ρ is not as transparent as that implied by our Inverse-Wishart prior. As discussed before, $\boldsymbol{W}_2^{-1}(3, \boldsymbol{I}_2)$ is a marginally uniform prior (Barnard, McCulloch and Meng 2000), which implies that $\rho \sim \text{Uniform}(-1, 1)$. Moreover, the reparametrization method is not directly applicable to the selection-Probit or selection-Robit, since the posterior distribution of ρ is nonconjugate or has a nonstandard form, which, however, can be easily solved by our parameter expansion scheme. The power of unification of all the selection models of parameter expansion motivates our Bayesian procedures.

In our first two empirical studies, frequentists' and Bayesian procedures give very similar results, which are consequences of the Bernstein-Von Mises theoreom (Van der Vaart, 2000) that the sampling distribution of the MLE and the posterior distribution for the same parameter have the same asymptotic Normality under regularity conditions and with large samples. However, they may provide different results when the regularity conditions fail or with small samples. In our third empirical study, the posterior distribution of ν is not unimodal, and the posterior distribution of the parameter ρ is very close to the boundary -1. The regularity conditions for asymptotic Normality may be violated, and inference based on Normal approximation may be inappropriate.

Acknowledgements

I want to thank Dr. Yulia V. Marchenko for providing him with the `Stata` command `heckt`, which implements the procedure in Marchenko and Genton (2012). And I am also grateful to the reviewers' constructive comments.

Appendices

Appendix A: Prior for $\boldsymbol{\Sigma}$ and Prior for $(\boldsymbol{\Omega}, \sigma_2^2)$

Let ω_{ij} denote the $(i,j)-$th element of $\boldsymbol{\Omega}$ with $\omega_{11} = \sigma_1^2$, and $\omega_{12} = \rho\sigma_1$. Let σ_{ij} denote the $(i,j)-$th element of $\boldsymbol{\Sigma}$ with $\sigma_{11} = \sigma_1^2 = \omega_{11}, \sigma_{12} = \rho\sigma_1\sigma_2 = \omega_{12}(\sigma_2^2)^{1/2}$ and $\sigma_{22} = \sigma_2^2$. The Jacobian matrix of the transformation $(\boldsymbol{\Omega}, \sigma_2^2) \to \boldsymbol{\Sigma}$ is

$$\frac{\partial(\sigma_{11}, \sigma_{12}, \sigma_{22})}{\partial(\omega_{11}, \omega_{12}, \sigma_2^2)} = \begin{array}{c} \\ \sigma_{11} \\ \sigma_{12} \\ \sigma_{22} \end{array} \begin{array}{c} \begin{array}{ccc} \omega_{11} & \omega_{12} & \sigma_2^2 \end{array} \\ \begin{pmatrix} 1 & 0 & 0 \\ 0 & \sigma_2 & \omega_{12}/(2\sigma_2) \\ 0 & 0 & 1 \end{pmatrix} \end{array},$$

and therefore the Jacobian of the transformation is $J\{(\boldsymbol{\Omega}, \sigma_2^2) \to \boldsymbol{\Sigma}\} = \sigma_2$.

If we use Inverse-Wishart prior $\boldsymbol{W}_2^{-1}(\nu_0, \boldsymbol{I}_2)$ for $\boldsymbol{\Sigma}$ with density

$$f(\boldsymbol{\Sigma}) \propto |\boldsymbol{\Sigma}|^{-(\nu_0+2+1)/2} \exp\left\{-\frac{1}{2}\text{tr}\left(\boldsymbol{\Sigma}^{-1}\right)\right\},$$

the prior for $(\sigma_2^2, \boldsymbol{\Omega})$ is

$$f(\sigma_2^2, \boldsymbol{\Omega}) \propto (\sigma_2^2)^{-\nu_0/2-1}|\boldsymbol{\Sigma}|^{-(\nu_0+2+1)/2} \exp\left\{-\frac{1}{2}\left(\omega^{11} + \frac{\omega^{22}}{\sigma_2^2}\right)\right\},$$

where ω^{ij} is the (i,j)-th element of $\boldsymbol{\Omega}^{-1}$. From joint distribution of $(\boldsymbol{\Omega}, \sigma_2^2)$, the prior for σ_2^2 given $\boldsymbol{\Omega}$ is $\sigma_2^2|\boldsymbol{\Omega} \sim \omega^{22}/\chi^2_{\nu_0} = \{(1-\rho^2)\chi^2_{\nu_0}\}^{-1}$, and the prior for $\boldsymbol{\Omega}$ is

$$\begin{aligned} f(\boldsymbol{\Omega}) &\propto |\boldsymbol{\Omega}|^{-(\nu_0+3)/2} e^{-\omega^{11}/2} \left(\omega^{22}\right)^{-\nu_0/2} \\ &\propto (1-\rho^2)^{-3/2}\sigma_1^{-(\nu_0+3)} \exp\left\{-\frac{1}{2\sigma_1^2(1-\rho^2)}\right\}. \end{aligned}$$

Appendix B: Prior for $\boldsymbol{\Sigma}$ and Prior for $(\sigma_1^2, \sigma_2^2, \boldsymbol{R})$

The Jacobian matrix of the transfomation $(\sigma_1^2, \sigma_2^2, \boldsymbol{R}) \to \boldsymbol{\Sigma}$ is

$$\frac{\partial(\sigma_{11}, \sigma_{12}, \sigma_{22})}{\partial(\sigma_1^2, \sigma_2^2, \rho)} = \begin{array}{c} \\ \sigma_{11} \\ \sigma_{12} \\ \sigma_{22} \end{array} \begin{array}{c} \begin{array}{ccc} \sigma_1^2 & \rho & \sigma_2^2 \end{array} \\ \begin{pmatrix} 1 & 0 & 0 \\ \rho(\sigma_2^2/\sigma_1^2)^{1/2} & \sigma_1\sigma_2 & \rho(\sigma_1^2/\sigma_2^2)^{1/2} \\ 0 & 0 & 1 \end{pmatrix} \end{array},$$

and therefore the Jacobian of the transformation is $J\{(\sigma_1^2, \sigma_2^2, \boldsymbol{R}) \to \boldsymbol{\Sigma}\} = \sigma_1\sigma_2$.

If we use Inverse-Wishart prior $\boldsymbol{W}_2^{-1}(\nu_0, \boldsymbol{I}_2)$ for $\boldsymbol{\Sigma}$, the prior for $(\sigma_1^2, \sigma_2^2, \boldsymbol{R})$ is

$$f(\sigma_1^2, \sigma_2^2, \boldsymbol{R}) \propto (\sigma_1\sigma_2)^{-(\nu_0+1)} |\boldsymbol{R}|^{-(\nu_0+2+1)} \exp\left\{-\frac{r^{11}}{2\sigma_1^2} - \frac{r^{22}}{2\sigma_2^2}\right\},$$

where r^{ij} is the (i,j)-th element of $\boldsymbol{R}^{-1}$. From the joint distribution of $(\sigma_1^2, \sigma_2^2, \boldsymbol{R})$, the priors for σ_1^2 and σ_2^2 are $\sigma_1^2|\boldsymbol{R} \sim r^{11}/\chi^2_{\nu_0} = \{(1-\rho^2)\chi^2_{\nu_0}\}^{-1}, \sigma_2^2|\boldsymbol{R} \sim r^{11}/\chi^2_{\nu_0} = \{(1-\rho^2)\chi^2_{\nu_0}\}^{-1}$, and they are independent given $\boldsymbol{R}$. The prior for $\boldsymbol{R}$ or equivalently ρ is $f(\rho) \propto (1-\rho^2)^{-(\nu_0-3)/2}$.

Appendix C: Gamma Approximation

Ignoring additive constants, the log conditional density of ν is

$$l(\nu) = N\nu\log(\nu/2)/2 - N\log\Gamma(\nu/2) + (\alpha_0 - 1)\log\nu - \xi\nu,$$

and the log density of Gamma(α^*, β^*) is

$$h(\nu) = (\alpha^* - 1)\log\nu - \beta^*\nu.$$

The first and second order derivatives of $l(\nu)$ and $h(\nu)$ are

$$l'(\nu) = \tfrac{N}{2}\log\left(\tfrac{\nu}{2}\right) + \tfrac{N}{2} - \tfrac{N}{2}\psi\left(\tfrac{\nu}{2}\right) + \tfrac{\alpha_0-1}{\nu} - \xi, \quad h'(\nu) = \frac{\alpha^*-1}{\nu} - \beta^*, \qquad \text{(A.1)}$$

$$l''(\nu) = \tfrac{N}{2\nu} - \tfrac{N}{4}\psi'\left(\tfrac{\nu}{2}\right) - \tfrac{\alpha_0-1}{\nu^2}, \qquad h''(\nu) = -\frac{\alpha^*-1}{\nu^2}, \qquad \text{(A.2)}$$

where $\psi(x) = \mathrm{d}\log\Gamma(x)/\mathrm{d}x$ is the digamma function, and $\psi'(x) = \mathrm{d}\psi(x)/\mathrm{d}x$ is the trigamma function. The mode of $h(\nu)$ is $(\alpha^*-1)/\beta^*$, and the curvature at the mode is $-\beta^{*2}/(\alpha^*-1)$. From (A.1) and (A.2), we can numerically find the mode of $l(\nu)$ denoted

as ν^*, and the curvature at the mode is $l^* = l''(\nu^*)$. By matching the modes and the curvatures at the modes of $l(\nu)$ and $h(\nu)$, the parameters of the Gamma approximation is chosen as

$$\alpha^* = 1 - \nu^{*2} l^* \text{ and } \beta^* = -\nu^* l^*. \tag{A.3}$$

References

Ahn, H. and Powell, J. L. (1993), "Semiparametric Estimation of Censored Selection Models with a Nonparametric Selection Mechanism," *Journal of Econometrics*, **58**, 3-29.

Barnard, J., McCulloch, R. and Meng, X. L. (2000), "Modeling covariance matrices in terms of standard deviations and correlations, with application to shrinkage," *Statistica Sinica*, **10**, 1281-1311.

Bärnighausen T., Bor J., Wandira-Kazibwe S., Canning D. (2011), "Correcting HIV prevalence estimates for survey nonparticipation using Heckman-type selection models," *Epidemiology*, **22**, 27-35.

Cameron, A. C. and Trivedi, P. K. (2010), *Microeconometrics Using Stata* (Revised Edition), College Station, TX: Stata Press.

Chib, S., Greenberg, E. and Jeliazkov, I. (2009), "Estimation of Semiparametric Models in the Presence of Endogeneity and Sample Selection," *Journal of Computational and Graphical Statistics*, **18**, 321-348.

Das, M., Newey, W. K. and Vella, F. (2003), "Nonparametric Estimation of Sample Selection Models," *Review of Economics Studies*, **70**, 33-58.

Geweke, J. (1992), "Priors for Macroeconomic Time Series and Their Application," Minneapolis: Federal Reserve Bank of Minneapolis, Institute for Empirical Macroeconomics discussion paper #64.

Heckman, J. J. (1979), "Sample Selection Bias as a Specification Error," *Econometrica*, **47**, 153-161.

Imai, K. and Van Dyk, D. A. (2005), "A Bayesian Analysis of the Multinomial Probit Model Using Marginal Data Augmentation," *Journal of Econometrics*, **124**, 311-

334.

Lee, L. F. (1983), "Generalized Econometric Models with Selectivity," *Econometrica*, **51**, 507-512.

Li, K. (1998), "Bayesian Inference in a Simultaneous Equation Model with Limited Dependent Variables," *Journal of Econometrics*, **85**, 387-400.

Little, R. J. and Rubin, D. B. (2002), *Statistical Analysis of Missing Data*, Wiley: New York.

Liu, C. (2004), "Robit Regression: A Simple Robust Alternative to Logistic and Probit Regression," In *Applied Bayesian Modeling and Causal Inference from Incomplete-Data Perspectives* (A. Gelman and X. L. Meng, eds.) 227-238. Wiley, New York.

Liu, C. (1999), "Efficient ML Estimation of the Multivariate Normal Distribution from Incomplete Data," *Journal of Multivariate Analysis*, **69**, 206-217.

Liu, C., Rubin, D. B. and Wu, Y. N. (1998), "Parameter Expansion to Accelerate EM: the PX-EM algorithm," *Biometrika*, **85**, 755-770.

Liu, J. S. (2001), *Monte Carlo Strategies in Scientific Computing*, NY: Springer-Verlag.

Liu, J. S. and Wu, Y. N. (1999), "Parameter Expansion for Data Augmentation," *Journal of the American Statistical Association*, **94**, 1264-1274.

Koop, G. and Poirier, D. J. (1997) Learning about the across-regime correlation in switching regression models. *Journal of Econometrics*, **78**, 217-227.

Marchenko, Y. V. and Genton, M. G. (2012), "A Heckman Selection-t Model," *Journal of the American Statistical Association*, **107**, 304-317.

McCulloch, R. E., Polson, N. G. and Rossi, P. E. (2000), "A Bayesian Analysis of the Multinomial Probit Model with Fully Identified Parameters," *Journal of Econometrics*, **99**, 173-193.

Meng, X. L. and Van Dyk, D. A. (1999), "Seeking Efficient Data Augmentation Schemes via Conditional and Marginal Augmentation," *Biometrika*, **86**, 301-320.

Mroz, T. A. (1987), "The Sensitivity of an Empirical Model of Married Womens Hours of Work to Economical and Statistical Assumptions," *Econometrica*, **55**, 765-799.

R Development Core Team (2010). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. ISBN

3-900051-07-0, URL http://www.R-project.org/.

StataCorp. (2013). Stata Statistical Software: Release 13. College Station, TX: StataCorp LP.

Tanner, M. A., and Wong, W. H. (1987), "The Calculation of Posterior Distributions by Data Augmentation" (with discussion), *Journal of the American Statistical Association*, **82**, 528-550.

Toomet, O. and Henningsen, A. (2008), "Sample Selection Models in R: Package sampleSelection," *Journal of Statistical Software*, **27**, http://www.jstatsoft.org/v27/i07.

Van der Vaart, A. W. (2000). *Asymptotic Statistics.* Cambridge University Press.

Van Dyk, D. A. and Meng, X. L. (2001), "The Art of Data Augmentation," *Journal of Computational and Graphical Statistics*, **10**, 1-50.

Van Hasselt, M. (2011), "Bayesian Inference in a Sample Selection Model," *Journal of Econometrics*, **165**, 221-232.

Wooldridge, J. M. (2002), *Econometric Analysis of Cross Section and Panel Data*, The MIT press.